

Scheduling for Optimal Rate Allocation in Ad Hoc Networks With Heterogeneous Delay Constraints

Juan José Jaramillo, *Member, IEEE*, R. Srikant, *Fellow, IEEE*, and Lei Ying, *Member, IEEE*

Abstract—This paper studies the problem of scheduling in single-hop wireless networks with real-time traffic, where every packet arrival has an associated deadline and a minimum fraction of packets must be transmitted before the end of the deadline. Using optimization and stochastic network theory we propose a framework to model the quality of service (QoS) requirements under delay constraints. The model allows for fairly general arrival models with heterogeneous constraints. The framework results in an optimal scheduling algorithm which fairly allocates data rates to all flows while meeting long-term delay demands. We also prove that under a simplified scenario our solution translates into a greedy strategy that makes optimal decisions with low complexity.

Index Terms—Wireless networks, ad hoc networks, quality of service, scheduling, real-time traffic.

I. INTRODUCTION

IN this paper we study the problem of scheduling real-time traffic in ad hoc networks under maximum per-packet delay constraints. The problem of scheduling best-effort traffic, which is defined as traffic that does not have any kind of quality of service (QoS) requirements such as minimum bandwidth or maximum delay, has been extensively studied for the case of wireless networks. An optimization framework for resource allocation in wireless networks has been developed in [1]–[6], where a dual decomposition approach was used to derive various components of the resource allocation architecture such as scheduling, congestion control, routing, power control, etc. A striking feature of the solution is an alternative derivation of the maxweight algorithm proposed in [7]. We refer the readers to [8], [9] for a survey of these works.

Scheduling algorithms for packets with strict deadline requirements have been proposed in [10]–[13], but the solutions are only approximate. In [14]–[16], the problem of optimal admission control and scheduling for real-time traffic was addressed for access-point wireless networks in which only one link can transmit at any given time. Among the many contributions in these papers is a key modeling innovation whereby the network is studied in frames, where a frame is a contiguous set of time-slots of fixed duration. Packets with deadlines are assumed to arrive only at the beginning of a

frame and have to be served before the end of the frame according to some specified deadlines.

The problem of optimal congestion control and scheduling for general ad hoc networks and arrivals was studied in [17], using the modeling paradigm of frames proposed in [14]. The model allows a common framework for handling both best-effort and real-time traffic simultaneously, but it can only handle homogeneous per-packet delay requirements.

In this paper, we further extend the results of [17] for the case of heterogeneous delays and more general arrival models, and study the impact of acknowledgments on scheduling in a lossy channel.

The main contributions of this work are summarized as follows:

- 1) We present an optimization formulation for the problem of scheduling real-time traffic under maximum per-packet delay constraints in wireless ad hoc networks. Unlike earlier work, the formulation allows for general arrival models with heterogeneous delays.
- 2) Using duality theory and a decomposition approach, we design an optimal scheduler that fairly allocates data rates to all links and ensures that a required fraction of each flow's packets are delivered before the prescribed deadline by appealing to connections between Lagrange multipliers and service deficits.
- 3) We prove that the scheduler meets all the QoS requirements and that converges to the optimal solution.
- 4) We then consider noisy channels where the transmitter does not have perfect channel state information and relies on feedback (acknowledgment) from the receiver to find out if a transmission was successful. We derive the structure of optimal scheduling algorithms in this case and show that in the special case of colocated networks, the scheduling algorithm reduces to a simple greedy algorithm. The last part recovers a key result in [14] using a different approach.

The paper is organized as follows. Section II presents the network model we use in this work. The optimization formulation is presented in Section III for the simplest of the channel models we study, while the dual decomposition approach is developed in Section IV. The optimal scheduler and its convergence properties are presented in Section V. Since we study three different channel models, in Sections VI and VII we highlight the differences between the simplest channel model and the other two, and show the relationship between feedback after every transmission and algorithm complexity. In Section VIII we perform a simulation study to understand

Research supported by NSF Grants 07-21286, 05-19691, 03-25673, 08-31756, 09-53165, ARO MURI Subcontracts, AFOSR Grant FA-9550-08-1-0432, and DTRA Grants HDTRA1-08-1-0016, HDTRA1-09-1-0055.

J. J. Jaramillo and L. Ying are with the Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011 USA. Email: {jjjarami, leiying}@iastate.edu

R. Srikant is with the Department of Electrical and Computer Engineering, and the Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801 USA. Email: rsrikant@illinois.edu

the rates that can be achieved under different channel models. Finally, in Section IX we present the conclusions.

II. NETWORK MODEL

In this section we present our model for a network composed of single-hop traffic flows, such that each packet has maximum delay requirements.

We represent the network using a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$, where \mathcal{N} is the set of nodes and \mathcal{L} is the set of links, such that for any $n_1, n_2 \in \mathcal{N}$, if $(n_1, n_2) \in \mathcal{L}$ then node n_1 can communicate to node n_2 . Links are numbered 1 through $|\mathcal{L}|$, and by abusing notation, we will sometimes use $l \in \mathcal{L}$ to mean $l \in \{1, 2, \dots, |\mathcal{L}|\}$.

Time is assumed to be divided in slots, and a set of T consecutive slots is called a *frame*. Let $\mathcal{T} \stackrel{\text{def}}{=} \{1, \dots, T\}$. We denote by $a = (a_{lt})_{l \in \mathcal{L}, t \in \mathcal{T}}$ the number of packet arrivals at a given frame for link l at time slot t , and assume that we get to know a at the beginning of the frame. Furthermore, assume that $a_l \stackrel{\text{def}}{=} \sum_{t \in \mathcal{T}} a_{lt}$ is a random variable with mean λ_l and variance σ_l^2 , such that $Pr(a_l = 0) > 0$ and $Pr(a_l = 1) > 0$. The last two assumptions are used to guarantee that the Markov chain to be defined later is both irreducible and aperiodic, but they can be replaced by similar assumptions. We further assume that arrivals are independent between different frames.

Define $\mathcal{T}_l^a \stackrel{\text{def}}{=} \{t : t \in \mathcal{T} \text{ and } a_{lt} > 0\}$ to be the set of arrival times at link l . Let $\tau = (\tau_{lt})_{l \in \mathcal{L}, t \in \mathcal{T}_l^a}$ be the deadline associated with packet arrivals. That is, a packet that arrived at link l , time t , must be transmitted by the end of time slot τ_{lt} . We assume that the deadlines are such that

$$\{t_1, \dots, \tau_{lt_1}\} \cap \{t_2, \dots, \tau_{lt_2}\} = \emptyset \text{ for all } t_1, t_2 \in \mathcal{T}_l^a$$

and

$$\tau_{lt} \leq T \text{ for all } l \in \mathcal{L}, t \in \mathcal{T}_l^a.$$

In other words, packets must be transmitted before the next set of arrivals occurs in subsequent time slots, and all packets must be transmitted before the end of the frame.

If a packet misses its deadline it is discarded, and it is required that the loss probability at link $l \in \mathcal{L}$ due to deadline expiry must be no more than p_l . To avoid unnecessary complexity in the formulation, we will write a to denote both the number of packet arrivals and its associated deadlines τ .

This paper studies the cases where the channel state is assumed to be constant for the duration of a frame and when it is allowed to change from time slot to time slot. In both cases we assume the state is independent between frames and independent of arrivals. When the channel is fixed in a frame, let $c = (c_l)_{l \in \mathcal{L}}$ denote the number of packets link l can successfully transmit in a time slot. When the channel is allowed to change, define $c = (c_{lt})_{l \in \mathcal{L}, t \in \mathcal{T}}$ to be the number of packets that can be successfully transmitted at link l in time slot t .

If we get to know the channel state before transmission, we can determine the optimal rate at which we can successfully transmit, possibly allowing more than one packet to be transmitted in a single slot. On the other hand, if the channel state is

not known, we can only determine whether a transmission was successful or not after we get some feedback from the receiver. In this paper we try to capture these different scenarios in the following cases:

- 1) Known channel state: It is assumed that c_l is a *non-negative* random variable with mean \bar{c}_l and variance $\sigma_{c_l}^2$, and we get to know the channel state at the *beginning* of the frame.
- 2) Unknown channel state, per-frame feedback: It is assumed that c_l is a *Bernoulli* random variable with mean \bar{c}_l and we only get to know the channel state at the *end* of the frame. In other words, the receiver acknowledges receptions for the all the packets in the frame, at the end of the frame.
- 3) Unknown channel state, per-slot feedback: It is assumed that c_{lt} is a *Bernoulli* random variable with mean \bar{c}_l and we get to know the channel state at the *end* of the time slot. In other words, acknowledgments are received after each transmission.

In the known channel state case we can potentially send more than one packet in a time slot at higher rates since channel estimation allow us to determine the optimal transmission rate. This is the reason why we make no assumptions on the values c_l can take since it will be determined by the particular wireless technology used. In the case where the channel is unknown before transmission we assume that we only get binary feedback in the form of acknowledgments, which is reflected in the Bernoulli assumption on c_l and c_{lt} . Thus, in this case, without any loss of generality, we assume only one packet can be transmitted per time slot per link.

For the sake of simplicity in the presentation, we will first develop the known channel case and we will later highlight the differences between this case and the other two cases in Sections VI and VII.

III. STATIC PROBLEM FORMULATION

To design our algorithm, we will first formulate the problem as a static optimization problem. Using duality theory, we will then obtain a dynamic solution to this problem and later we will prove its properties using stochastic Lyapunov techniques.

Let $s = (s_{lt})_{l \in \mathcal{L}, t \in \mathcal{T}}$ denote the number of packets scheduled for transmission at link l and time slot t . We will only focus on feasible schedules, so if $s_{l_1 t} > 0$ and $s_{l_2 t} > 0$ for any t , then links l_1 and l_2 can be scheduled to simultaneously transmit without interfering with each other.

Since we cannot transmit more packets than what are available and what the channel state allows, we have the following constraints when the arrivals are given by a and the channel state is c :

$$\sum_{j=t}^{\tau_{lt}} s_{lj} \leq a_{lt} \text{ for all } t \in \mathcal{T}_l^a, l \in \mathcal{L}, \quad (1)$$

$$s_{lt} = 0 \text{ for all } t \in \mathcal{T} \setminus \bigcup_{l \in \mathcal{L}} \mathcal{T}_l^a, l \in \mathcal{L}, \text{ and } \quad (2)$$

$$s_{lt} \leq c_l \text{ for all } l \in \mathcal{L} \text{ and } t \in \mathcal{T}. \quad (3)$$

Denote the set of feasible schedules when the arrivals and channel state are a and c by $\mathcal{S}(a, c)$, capturing any interference

constraints imposed by the network and satisfying (1), (2), and (3).

Our goal is to find $Pr(s|a, c)$ which is the probability of using schedule $s \in \mathcal{S}(a, c)$ when the arrivals are given by a and the channel state is c , subject to the constraint that the loss probability at link $l \in \mathcal{L}$ due to deadline expiry cannot exceed p_l .

Denoting by $\mu(a, c) = (\mu_l(a, c))_{l \in \mathcal{L}}$ the expected number of packets served when the arrivals and channel state are given by a and c , respectively, we have:

$$\mu_l(a, c) \leq \sum_{s \in \mathcal{S}(a, c)} \sum_{t \in \mathcal{T}} s_{lt} Pr(s|a, c) \text{ for all } l \in \mathcal{L}.$$

Thus, the expected service at link $l \in \mathcal{L}$ is given by

$$\mu_l \stackrel{def}{=} \sum_{a, c} \mu_l(a, c) Pr(c) Pr(a).$$

Due to QoS constraints we need at all links

$$\mu_l \geq \lambda_l(1 - p_l),$$

and to avoid trivialities, we assume that $\lambda_l(1 - p_l) > 0$ for all $l \in \mathcal{L}$.

For notational simplicity, define the capacity region for fixed arrivals and channel state as

$$\mathcal{C}(a, c) \stackrel{def}{=} \left\{ (\bar{\mu}_l)_{l \in \mathcal{L}} : \text{there exists } \bar{s} \in \mathcal{S}(a, c)_{c\mathcal{H}}, \right. \\ \left. \bar{\mu}_l \leq \sum_{t \in \mathcal{T}} \bar{s}_{lt}, \right\},$$

where $\mathcal{S}(a, c)_{c\mathcal{H}}$ is the convex hull of $\mathcal{S}(a, c)$.

Thus, the overall capacity region can be defined as follows:

$$\mathcal{C} \stackrel{def}{=} \left\{ (\mu_l)_{l \in \mathcal{L}} : \text{there exists } (\bar{\mu}_l(a, c))_{l \in \mathcal{L}} \in \mathcal{C}(a, c) \right. \\ \left. \text{for all } a, c \text{ and } \mu_l = E[\bar{\mu}_l(a, c)] \text{ for all } l \in \mathcal{L} \right\}.$$

We will focus on the following static formulation for our problem, for some given vector $w \in \mathbb{R}_+^{|\mathcal{L}|}$:

$$\max_{\mu \in \mathcal{C}} \sum_{l \in \mathcal{L}} w_l \mu_l \quad (4)$$

subject to

$$\mu_l \geq \lambda_l(1 - p_l) \text{ for all } l \in \mathcal{L}.$$

The vector w can be used to allocate additional bandwidth fairly to flows beyond what is required to meet their QoS needs. Other uses for w have been explored in [17]. We will assume that the arrivals and loss probability requirements are feasible and thus the optimization problem has a solution μ^* .

IV. DUAL DECOMPOSITION OF THE STATIC PROBLEM

In this section we use duality theory to decompose the static optimization problem into simpler subproblems that will give us the ideas behind the dynamic algorithm.

Using the definition of the dual function [18], we have that

$$D(\delta) = \max_{\mu \in \mathcal{C}} \sum_{l \in \mathcal{L}} w_l \mu_l - \delta_l [\lambda_l(1 - p_l) - \mu_l].$$

From Slater's condition [19] we know that the duality gap is zero and therefore $D(\delta^*) = \sum_{l \in \mathcal{L}} w_l \mu_l^*$, where

$$\delta^* \in \arg \min_{\delta_l \geq 0} D(\delta).$$

We are interested in finding μ^* but not the value $D(\delta^*)$, so the problem can be simplified as follows

$$\max_{\mu \in \mathcal{C}} \sum_{l \in \mathcal{L}} (w_l + \delta_l) \mu_l. \quad (5)$$

Since we are interested in solving the problem for non-negative values of δ_l , it must be the case that μ^* is as large as the constraints allow. Furthermore, since the objective function in (5) is linear, the problem can be decomposed into the following subproblems for fixed a and c :

$$\max_{s \in \mathcal{S}(a, c)} \sum_{l \in \mathcal{L}} (w_l + \delta_l) \sum_{t \in \mathcal{T}} s_{lt}.$$

The optimization problem can be solved using the following iterative algorithm, where k is the step index:

$$\tilde{s}^*(a, c, k) \in \arg \max_{s \in \mathcal{S}(a, c)} \sum_{l \in \mathcal{L}} [w_l + \delta_l(k)] \sum_{t \in \mathcal{T}} s_{lt}$$

$$\tilde{\mu}_l^*(k) = \sum_{a, c} \sum_{t \in \mathcal{T}} \tilde{s}_{lt}^*(a, c, k) Pr(c) Pr(a).$$

And the update equation for the Lagrange multipliers is given by

$$\delta_l(k+1) = \{\delta_l(k) + \epsilon[\lambda_l(1 - p_l) - \tilde{\mu}_l^*(k)]\}^+,$$

where $\epsilon > 0$ is a fixed step-size parameter, and for any $\alpha \in \mathbb{R}$, $\alpha^+ \stackrel{def}{=} \max\{\alpha, 0\}$.

With the change of variables $\epsilon \hat{d}(k) = \delta(k)$, we rewrite the algorithm as

$$\tilde{s}^*(a, c, k) \in \arg \max_{s \in \mathcal{S}(a, c)} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + \hat{d}_l(k) \right] \sum_{t \in \mathcal{T}} s_{lt} \quad (6)$$

$$\tilde{\mu}_l^*(k) = \sum_{a, c} \sum_{t \in \mathcal{T}} \tilde{s}_{lt}^*(a, c, k) Pr(c) Pr(a),$$

with update equation:

$$\hat{d}_l(k+1) = [\hat{d}_l(k) + \lambda_l(1 - p_l) - \tilde{\mu}_l^*(k)]^+.$$

From the update equation we see that $\hat{d}_l(k)$ can be interpreted as a queue with $\lambda_l(1 - p_l)$ arrivals and $\tilde{\mu}_l^*(k)$ departures at step k .

V. ONLINE ALGORITHM AND ITS CONVERGENCE ANALYSIS

So far we have presented a dual decomposition for a static problem; however, the real network has stochastic arrivals and channel state. We will use the intuition from the decomposition in this section to develop an online algorithm that can cope with such changing conditions and prove its convergence properties.

A. Scheduler

We propose the following dynamic scheduling algorithm, where the arrivals and channel state in frame k are given by $a(k)$ and $c(k)$, respectively:

$$\tilde{s}^*(a(k), c(k), d(k)) \in \arg \max_{s \in \mathcal{S}(a(k), c(k))} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + d_l(k) \right] \sum_{t \in \mathcal{T}} s_{lt}, \quad (7)$$

with update equation

$$d_l(k+1) = [d_l(k) + \tilde{a}_l(k) - I_l^*(a(k), c(k), d(k))]^+,$$

where

$$I_l^*(a(k), c(k), d(k)) = \sum_{t \in \mathcal{T}} \tilde{s}_{lt}^*(a(k), c(k), d(k))$$

and $\tilde{a}_l(k)$ is a binomial random variable with parameters $a_l(k)$ and $1 - p_l$. The quantity $\tilde{a}_l(k)$ can be generated by the network as follows: upon each packet arrival, toss a coin with probability of *heads* equal to $1 - p_l$, and if the outcome is *heads*, add a one to the deficit counter $d_l(k)$. This implementation for $\tilde{a}_l(k)$ was first suggested in [17].

Note that in (7) we make explicit the fact that the optimal scheduler is a function of $a(k)$, $c(k)$, and $d(k)$, for fixed w and ϵ . Also note that $d_l(k)$ can be interpreted as a virtual queue that keeps track of the deficit in service for link l to achieve a loss probability due to deadline expiry less than or equal to p_l . The idea of using a deficit counter was first used in [14] for the case of colocated networks with homogeneous delays, while the Lagrange multiplier interpretation allowed us to extend the result to general ad hoc networks and heterogeneous delays.

B. Convergence Results

We now prove that the online algorithm meets the QoS constraints, the total expected service deficit has a $O(1/\epsilon)$ bound, and the expected value of the objective is within $O(\epsilon)$ of the optimal value of the static problem (4). For the sake of readability, we defer the proofs to the appendices.

We will first bound the expected drift of $d(k)$ for a suitable Lyapunov function. Note that $d(k)$ defines an irreducible and aperiodic Markov chain.

Lemma 1: Consider the Lyapunov function $V(d) = \frac{1}{2} \sum_{l \in \mathcal{L}} d_l^2$. If there exists a point $\mu(\Delta) \in \mathcal{C}/(1 + \Delta)$ for some $\Delta > 0$ such that

$$\mu_l(\Delta) \geq \lambda_l(1 - p_l) \text{ for all } l \in \mathcal{L} \quad (8)$$

then

$$\begin{aligned} E[V(d(k+1)) | d(k) = d] - V(d) &\leq B_1 - B_2 \sum_{l \in \mathcal{L}} d_l \\ &\quad - \frac{1}{\epsilon} \sum_{l \in \mathcal{L}} \{w_l(1 + \Delta)\mu_l(\Delta) - w_l E[I_l^*(a(k), c(k), d)]\} \end{aligned}$$

for some positive constants B_1 , B_2 , any $\epsilon > 0$, where $I^*(a(k), c(k), d)$ is obtained from the solution to (7). \diamond

Since the last term in the right-hand side of the inequality can be upper-bounded, it can be shown that the expected drift is negative but for a finite set of values of $d(k)$. Thus, Lemma

1 implies that $d(k)$ is positive recurrent. As a corollary of this result, we have that the total expected service deficit has an $O(1/\epsilon)$ bound.

Corollary 1: If there exists a point $\mu(\Delta) \in \mathcal{C}/(1 + \Delta)$ for some $\Delta > 0$ such that (8) holds true, then the total expected service deficit is upper-bounded by

$$\limsup_{k \rightarrow \infty} E \left[\sum_{l \in \mathcal{L}} d_l(k) \right] \leq B_3 + \frac{1}{\epsilon} B_4$$

for all $l \in \mathcal{L}$, where $B_3 = B_1/B_2$ and $B_4 = \sum_{l \in \mathcal{L}} w_l \lambda_l / B_2$. \diamond

Lemma 1 also implies that the scheduler fulfills all QoS requirements.

Corollary 2: If there exists a point $\mu(\Delta) \in \mathcal{C}/(1 + \Delta)$ for some $\Delta > 0$ such that (8) holds true, then the online algorithm fulfills all the QoS constraints. That is:

$$\liminf_{K \rightarrow \infty} E \left[\frac{1}{K} \sum_{k=1}^K I_l^*(a(k), c(k), d(k)) \right] \geq \lambda_l(1 - p_l)$$

for all $l \in \mathcal{L}$. \diamond

The above corollary states that the arrival rate into the deficit counter is less than or equal to the service rate. The result is an obvious consequence of the stability of the deficit counters and so a formal proof is not given.

In order to be able to prove that our scheduler achieves the optimal solution to the static problem in an average sense, a related result to Lemma 1 must be stated first. Since the proof is similar to the proof in Lemma 1, it is omitted.

Lemma 2: Consider the Lyapunov function $V(d) = \frac{1}{2} \sum_{l \in \mathcal{L}} d_l^2$. Then

$$\begin{aligned} E[V(d(k+1)) | d(k) = d] - V(d) &\leq B_1 - B_5 \sum_{l \in \mathcal{L}} d_l \\ &\quad - \frac{1}{\epsilon} \sum_{l \in \mathcal{L}} \{w_l \mu_l^* - w_l E[I_l^*(a(k), c(k), d)]\} \end{aligned}$$

for $B_1 > 0$, some non-negative constant B_5 , any $\epsilon > 0$, where μ^* is a solution to (4), and $I^*(a(k), c(k), d)$ is obtained from the solution to (7). \diamond

Using Lemma 2 we can now prove that our online algorithm is within $O(\epsilon)$ of the optimal value.

Theorem 1: For any $\epsilon > 0$ we have that

$$\limsup_{K \rightarrow \infty} E \left[\sum_{l \in \mathcal{L}} w_l \mu_l^* - w_l \frac{1}{K} \sum_{k=1}^K I_l^*(a(k), c(k), d(k)) \right] \leq B\epsilon$$

for some $B > 0$, where μ^* is a solution to (4), and $I^*(a(k), c(k), d(k))$ is obtained from the solution to (7). \diamond

Both Corollary 1 and Theorem 1 highlight the trade-off in choosing ϵ : the closer we get to the optimal value for the static formulation, the more the aggregate in the deficit counters will increase.

The statement and proofs of Lemma 1 and Theorem 1 follow the techniques in [17], which are similar to the techniques in [3]. Slightly different results can be derived using the techniques in [4] and [1].

VI. UNKNOWN CHANNEL STATE, PER-FRAME FEEDBACK

Since the analysis for the unknown channel state, per-frame feedback case is similar in nature to the known channel case, in this section we will only highlight the differences between both cases.

Remember that we assume that the channel at link $l \in \mathcal{L}$, c_l , is a *Bernoulli* random variable with mean \bar{c}_l , and we get to know the channel state only at the end of the frame. Thus, the set of feasible schedules when the arrivals are given by a is the set of schedules that capture any interference constraints imposed by the network and that fulfill the following constraints:

$$\sum_{j=t}^{\tau_{lt}} s_{lj} \leq a_{lt} \text{ for all } t \in \mathcal{T}_l^a, l \in \mathcal{L}, \quad (9)$$

$$s_{lt} = 0 \text{ for all } t \in \mathcal{T} \setminus \cup_{l \in \mathcal{L}} \{t, \dots, \tau_{lt}\}, l \in \mathcal{L}, \text{ and } \quad (10)$$

$$s_{lt} \leq 1 \text{ for all } l \in \mathcal{L} \text{ and } t \in \mathcal{T}. \quad (11)$$

From our assumption that the channel state remains constant for the duration of a frame, (9) and (10) tell us that we should not schedule a link for transmission more than the number of packets that are available, since extra transmissions do not increase the number of successes. Furthermore, since the channel is Bernoulli, when a link is scheduled it can only transmit at most a packet, as the constraint in (11) indicates. Let $\mathcal{S}(a)$ denote the set of feasible schedules when the arrivals are given by a , capturing any interference constraints given by the network and fulfilling (9), (10), and (11).

Following the same arguments as in Section III, we have to design a scheduling strategy $Pr(s|a)$ that is the probability of using schedule $s \in \mathcal{S}(a)$ when the arrivals are given by a . Since we cannot base our policy on the state of the channel, as we did before, the expected service to link $l \in \mathcal{L}$ is now given by

$$\mu_l \leq \sum_{a,c} \sum_{s \in \mathcal{S}(a)} \sum_{t \in \mathcal{T}} c_l s_{lt} Pr(s|a) Pr(c) Pr(a).$$

Simplifying the equation we get

$$\mu_l \leq \sum_a \sum_{s \in \mathcal{S}(a)} \sum_{t \in \mathcal{T}} \bar{c}_l s_{lt} Pr(s|a) Pr(a). \quad (12)$$

That is, the expected service depends on the average channel state. Properly modifying the definition of the capacity of the network, we can write the static formulation as in (4).

Using the dual decomposition technique of Section IV we can develop the design ideas behind the following dynamic scheduler, assuming that at frame k the arrivals are given by $a(k)$:

$$\tilde{s}^*(a(k), d(k)) \in \arg \max_{s \in \mathcal{S}(a(k))} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + d_l(k) \right] \bar{c}_l \sum_{t \in \mathcal{T}} s_{lt}, \quad (13)$$

with update equation

$$d_l(k+1) = [d_l(k) + \tilde{a}_l(k) - I_l^*(a(k), d(k))]^+,$$

where

$$I_l^*(a(k), d(k)) = \sum_{t \in \mathcal{T}} \tilde{s}_{lt}^*(a(k), d(k))$$

and $\tilde{a}_l(k)$ is a binomial random variable with parameters $a_l(k)$ and $1 - p_l$.

The main difference in the scheduler compared to the known channel case is that the algorithm now uses the expected channel state in making scheduling decisions. Thus, the algorithm needs to know or estimate \bar{c}_l .

Using the same techniques as in Section V-B, we can prove that the scheduler meets all the QoS requirements, the total expected service deficits have an $O(1/\epsilon)$ bound, and the mean value of the objective is within $O(\epsilon)$ of the optimal value.

VII. UNKNOWN CHANNEL STATE, PER-SLOT FEEDBACK

Compared to the previous two cases, the per-slot feedback case is more complex due to the fact that we can use the feedback to update our decisions at every time slot. In this section we will first formulate the problem focusing on policies rather than on schedules, and we will show that no simple decomposition can be achieved for this case. However, we will prove that for a simple scenario a greedy solution can achieve the optimal solution.

A. Problem Formulation and Solution

Similar to the development in Section VI, we will only highlight the differences between this case and the known channel state case.

As described in Section II, the channel at link $l \in \mathcal{L}$, time slot $t \in \mathcal{T}$, c_{lt} , is assumed to be a Bernoulli random variable with mean \bar{c}_l . Thus, instead of choosing a schedule for the entire frame, we will try to find a scheduling policy ρ that makes decisions at every time slot based on the feedback received. Note however that if the arrivals and channel state at a given frame are given by a and c respectively, policy ρ will generate a schedule by the end of the frame. We will denote by $s(\rho, a, c)$ such schedule.

We only focus on feasible policies, which are defined to be policies that generate a schedule that meets all interference constraints given by the network and fulfill the following constraints, for fixed ρ , a , and c :

$$s_{tt}(\rho, a, c) = 0 \text{ for all } t \in \mathcal{T} \setminus \cup_{l \in \mathcal{L}} \{t, \dots, \tau_{lt}\}, l \in \mathcal{L}, \quad (14)$$

$$s_{lt}(\rho, a, c) \leq 1 \text{ for all } l \in \mathcal{L} \text{ and } t \in \mathcal{T}, \text{ and } \quad (15)$$

$$\sum_{j=t}^{\tau_{lt}} c_{lj} s_{lj}(\rho, a, c) \leq a_{lt} \text{ for all } t \in \mathcal{T}_l^a, l \in \mathcal{L}. \quad (16)$$

Note that (14) specifies that a link should not be scheduled if there is no packet to be transmitted, (15) states that we cannot schedule more than a packet in a time slot since the channel is Bernoulli, and (16) specifies that a feasible policy cannot have more successful transmissions than the number of packets available.

We highlight the fact that since there is only a finite number of feasible schedules, then the set of feasible policies is finite. We will denote by $\mathcal{P}(a)$ the set of feasible policies that meet all interference constraints and that fulfill (14), (15), and (16), when arrivals are given by a .

Our goal is to find the probability distribution $Pr(\rho|a)$ of using policy $\rho \in \mathcal{P}(a)$ in a given frame when arrivals are given

by a , such that the fraction of packets that miss the deadline at link l cannot exceed p_l . Thus, the expected service at link $l \in \mathcal{L}$ is subject to the following constraint

$$\mu_l \leq \sum_{a,c} \sum_{\rho \in \mathcal{P}(a)} \sum_{t \in \mathcal{T}} c_{lt} s_{lt}(\rho, a, c) Pr(\rho|a) Pr(c) Pr(a).$$

Therefore, the optimization problem is as follows, for a given vector $w \in \mathbb{R}_{+}^{|\mathcal{L}|}$:

$$\max_{\mu, Pr(\rho|a)} \sum_{l \in \mathcal{L}} w_l \mu_l \quad (17)$$

subject to

$$\mu_l \leq \sum_{a,c} \sum_{\rho \in \mathcal{P}(a)} \sum_{t \in \mathcal{T}} c_{lt} s_{lt}(\rho, a, c) Pr(\rho|a) Pr(c) Pr(a) \text{ for all } l$$

$$\mu_l \geq \lambda_l(1 - p_l) \text{ for all } l \in \mathcal{L}$$

$$Pr(\rho|a) \geq 0 \text{ for all } \rho \in \mathcal{P}(a), a$$

$$\sum_{\rho \in \mathcal{P}(a)} Pr(\rho|a) \leq 1 \text{ for all } a.$$

We will assume that the arrivals and loss probability requirements are feasible and thus the optimization problem has a solution μ^* .

Following the arguments in Section IV, we can develop the design ideas behind the following dynamic scheduler, assuming that at frame k the arrivals are given by $a(k)$ and the channel state by $c(k)$:

$$\tilde{\rho}^*(a(k), d(k)) \in \arg \max_{\rho \in \mathcal{P}(a(k))} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + d_l(k) \right] \mu_l(\rho, a(k)) \quad (18)$$

with update equation

$$d_l(k+1) = [d_l(k) + \tilde{a}_l(k) - I_l^*(a(k), c(k), d(k))]^+,$$

where

$$\mu_l(\rho, a(k)) = \sum_c \sum_{t \in \mathcal{T}} c_{lt} s_{lt}(\rho, a(k), c) Pr(c), \quad (19)$$

$$I_l^*(a(k), c(k), d(k)) = \sum_{t \in \mathcal{T}} c_{lt}(k) s_{lt}(\tilde{\rho}^*(k), a(k), c(k)),$$

$\tilde{\rho}^*(k) = \tilde{\rho}^*(a(k), d(k))$, and $\tilde{a}_l(k)$ is a binomial random variable with parameters $a_l(k)$ and $1 - p_l$.

We note that compared to the known channel case, the algorithm needs to know the probability distribution of c in order to make optimal decisions. From (18) we see that the duality approach does not give us a simple decomposition as in (6). The reason comes from the fact that even though per-slot feedback may help us to potentially increase the throughput, it also increases the complexity of the decision algorithm.

Using the same proof techniques as in Section V-B, it can be proved that the scheduler meets all the QoS requirements, the total expected service deficits have an $O(1/\epsilon)$ bound, and the mean value of the objective is within $O(\epsilon)$ of the optimal value.

B. A Greedy Strategy for Colocated Networks

In this Section we will show that in a simple scenario a greedy algorithm can achieve the optimal solution with minimum complexity. To do that, we will focus our attention to colocated networks, where only one link is allowed to transmit at any given time slot. We will also assume that the channel state is independent between different time slots. Furthermore, we will assume that at every frame there is a single packet arrival at every link at the beginning of the frame, and that all the packets must be transmitted by the end of the frame. That is,

$$\mathcal{T}_l^a = \{1\}, a_{l1} = 1, \text{ and } \tau_{l1} = T \text{ for all } l \in \mathcal{L}. \quad (20)$$

The key idea we will use in this section is that for a given frame when the deficit counters are given by d , links will be prioritized in decreasing order of the priorities $[\frac{1}{\epsilon} w_l + d_l] \bar{c}_l$.

Definition 1: A *greedy policy* for colocated networks is a scheduling policy that at every time slot schedules a link with the highest priority $[\frac{1}{\epsilon} w_l + d_l] \bar{c}_l$ among the links that have a packet that remains to be transmitted. \diamond

Theorem 2: The greedy scheduler is the optimal solution to (18) for colocated networks, when the arrivals are given by (20), and the channel state is independent between different time slots. \diamond

For ease of readability, we defer the proof to the appendixes. Due to the optimality of the policy, and following a similar development as in Section V-B, one can prove that the greedy algorithm meets all the QoS constraints, the total expected service deficits have an $O(1/\epsilon)$ bound, and the mean value of the objective is within $O(\epsilon)$ of the optimal value. We skip the proofs since they are analogous to the ones already presented.

The above theorem shows that the dual decomposition solution presented here recovers the solution for the special case of access-point networks presented in [14]. The contribution of this section is to show that the dual approach allow us to extend such results for very general ad hoc networks, arrivals, and for heterogeneous delays, and that [14] can be seen as a particular case of our general formulation.

VIII. SIMULATIONS

The purpose of the simulations is to compare the throughput that can be achieved for the different channel models. To do that, we simulate a 10-link network such that its interference graph is given by Fig. 1, where each vertex of the graph represents a link and the edges represent the interference constraints. For example, if link 1 is scheduled, then links 2, 4, and 7 cannot be activated. This interference graph was first used in [17].

The constraint for the dropping probability due to deadline expiration is set to 0.1, the packet arrivals at every link are assumed to be Bernoulli random variables with mean 0.6 packets/frame, and a frame has 3 time slots. Every channel is assumed to be a Bernoulli random variable with mean 0.96. We will compare the three channel models studied in this paper: known channel state, unknown channel state per-frame feedback, and unknown channel state per-slot feedback. The simulation time is 10^6 frames.




Fig. 1. Interference graph used in the simulations




Fig. 2. Average service when $w_l = 0$

In Figs. 2 and 3 we plot the average service for different values of w_l . The case $w_l = 0$ means that we are only interested in finding a feasible solution without any concern of optimality. This results in a rate assignment that is only slightly above the minimum required to achieve an acceptable dropping probability, as shown in Fig. 4. In the case $w_l = 6$ we have that the algorithm tries to maximize the service rate for all links, resulting in a significant decrease of the dropping probability as shown in Fig. 5. This result suggests that the objective function not only has a role giving priorities to links according to its weights, but also has a role in decreasing the dropping probabilities, allocating any available capacity to links.

It is interesting to note that the difference in service rates for the different channel models is smaller than 2%, which suggests that the added complexity in solving the optimal scheduler for the per-slot feedback case does not compare with the gains that can be achieved in terms of throughput. Thus, the per-frame feedback case can be used as a low-complexity approximation to the per-slot feedback case.

IX. CONCLUSIONS

In this work we have presented an optimization formulation for the problem of scheduling real-time traffic in ad hoc networks under maximum delay constraints. The model allows for general arrival models with heterogeneous delay constraints. Using duality theory and a decomposition approach, we presented an optimal scheduler and proved that it fairly allocates data rates to all links and guarantees that the delay requirements are met at every flow. We further studied the impact of feedback at every time slot on the complexity of




Fig. 3. Average service when $w_l = 6$




Fig. 4. Dropping probability when $w_l = 0$

the optimal algorithm, and showed that for a certain simple scenario a greedy strategy can achieve the optimal solution with low complexity, recovering the results of [14] for access-point networks.

APPENDIX A PROOF OF LEMMA 1

Before we prove Lemma 1, we present the following fact.

Fact 1: The optimization in (7) can be performed over $\mathcal{S}(a(k), c(k))_{\mathcal{CH}}$, the convex hull of $\mathcal{S}(a(k), c(k))$; that is,

$$\begin{aligned} \arg \max_{s \in \mathcal{S}(a(k), c(k))_{\mathcal{CH}}} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + d_l(k) \right] \sum_{t \in \mathcal{T}} s_{lt} = \\ \arg \max_{s \in \mathcal{S}(a(k), c(k))_{\mathcal{CH}}} \sum_{l \in \mathcal{L}} \left[\frac{1}{\epsilon} w_l + d_l(k) \right] \sum_{t \in \mathcal{T}} s_{lt}. \end{aligned}$$

The reason for this comes from the fact that the objective function is linear and therefore there must be an optimal point $\tilde{s}^*(a(k), c(k), d(k)) \in \mathcal{S}(a(k), c(k))$. \diamond




Fig. 5. Dropping probability when $w_l = 6$

Proof of Lemma 1:

$$\begin{aligned}
& E[V(d(k+1))|d(k)=d] - V(d) \\
&= E\left[\frac{1}{2} \sum_{l \in \mathcal{L}} \{[d_l + \tilde{a}_l(k) - I_l^*(a(k), c(k), d)]^+\}^2\right] - \sum_{l \in \mathcal{L}} \frac{d_l^2}{2} \\
&\leq E\left[\frac{1}{2} \sum_{l \in \mathcal{L}} [d_l + \tilde{a}_l(k) - I_l^*(a(k), c(k), d)]^2\right] - \sum_{l \in \mathcal{L}} \frac{d_l^2}{2} \\
&= E\left[\sum_{l \in \mathcal{L}} d_l [\tilde{a}_l(k) - I_l^*(a(k), c(k), d)] \right. \\
&\quad \left. + \frac{1}{2} \sum_{l \in \mathcal{L}} [\tilde{a}_l(k) - I_l^*(a(k), c(k), d)]^2\right] \\
&\leq E\left[\sum_{l \in \mathcal{L}} d_l \tilde{a}_l(k) - d_l I_l^*(a(k), c(k), d) \right. \\
&\quad \left. + \frac{1}{2} \sum_{l \in \mathcal{L}} \tilde{a}_l^2(k) + a_l^2(k)\right] \\
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \lambda_l (1 - p_l) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) I_l^*(a(k), c(k), d) \right. \\
&\quad \left. - \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l I_l^*(a(k), c(k), d)\right] \\
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) I_l^*(a(k), c(k), d) \right. \\
&\quad \left. - \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l I_l^*(a(k), c(k), d)\right] \\
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) I_l^*(a(k), c(k), d) \right. \\
&\quad \left. - \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l I_l^*(a(k), c(k), d)\right]
\end{aligned} \tag{21}$$

$$\begin{aligned}
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) I_l^*(a(k), c(k), d) \right. \\
&\quad \left. - \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l I_l^*(a(k), c(k), d)\right]
\end{aligned} \tag{22}$$

where (21) follows from the definition of $I_l^*(a(k), c(k), d)$, (22) follows from (8), and

$$B_1 = \frac{1}{2} \sum_{l \in \mathcal{L}} (\lambda_l^2 + \sigma_l^2) [1 + (1 - p_l)^2] + \lambda_l p_l (1 - p_l).$$

Given the definition of \mathcal{C} , we have that $\mu(\Delta) \in \mathcal{C}/(1 + \Delta)$ implies that there exist $\bar{\mu}(a, c) \in \mathcal{C}(a, c)$ for all a, c and $(1 + \Delta)\mu_l(\Delta) = E[\bar{\mu}_l(a, c)]$ for all $l \in \mathcal{L}$. For the rest of the proof we define $\bar{\mu}(a(k), c(k))$ to be such set of values associated to $(1 + \Delta)\mu(\Delta)$. Hence:

$$\begin{aligned}
& E[V(d(k+1))|d(k)=d] - V(d) \\
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) I_l^*(a(k), c(k), d)\right] \\
&\quad + \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l E[I_l^*(a(k), c(k), d)] \\
&\leq B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - E\left[\sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) \bar{\mu}_l(a(k), c(k))\right] \\
&\quad + \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l E[I_l^*(a(k), c(k), d)] \\
&= B_1 + \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) - \sum_{l \in \mathcal{L}} \left(\frac{1}{\epsilon} w_l + d_l\right) (1 + \Delta) \mu_l(\Delta) \\
&\quad + \sum_{l \in \mathcal{L}} \frac{1}{\epsilon} w_l E[I_l^*(a(k), c(k), d)] \\
&= B_1 - \Delta \sum_{l \in \mathcal{L}} d_l \mu_l(\Delta) \\
&\quad - \frac{1}{\epsilon} \sum_{l \in \mathcal{L}} \{w_l (1 + \Delta) \mu_l(\Delta) - w_l E[I_l^*(a(k), c(k), d)]\},
\end{aligned} \tag{23}$$

where (23) follows from the fact that $\bar{\mu}(a(k), c(k)) \in \mathcal{C}(a(k), c(k))$ and Fact 1. Therefore:

$$\begin{aligned}
& E[V(d(k+1))|d(k)=d] - V(d) \leq B_1 - B_2 \sum_{l \in \mathcal{L}} d_l \\
&\quad - \frac{1}{\epsilon} \sum_{l \in \mathcal{L}} \{w_l (1 + \Delta) \mu_l(\Delta) - w_l E[I_l^*(a(k), c(k), d)]\}
\end{aligned}$$

where

$$B_2 = \Delta \min_{l \in \mathcal{L}} \{\mu_l(\Delta)\}.$$

■

APPENDIX B PROOF OF THEOREM 1

From Lemma 2 we know that

$$\begin{aligned}
& \frac{1}{\epsilon} \sum_{l \in \mathcal{L}} \{w_l \mu_l^* - w_l E[I_l^*(a(k), c(k), d)]\} \\
&\leq B_1 - B_5 \sum_{l \in \mathcal{L}} d_l + V(d) - E[V(d(k+1))|d(k)=d]
\end{aligned}$$

$$\leq B_1 + V(d) - E[V(d(k+1)) | d(k) = d]$$

since $B_5 \sum_{l \in \mathcal{L}} d_l \geq 0$. Taking expectations:

$$\begin{aligned} & \frac{1}{\epsilon} E \left[\sum_{l \in \mathcal{L}} \{w_l \mu_l^* - w_l I_l^*(a(k), c(k), d(k))\} \right] \\ & \leq B_1 - E[V(d(k+1))] + E[V(d(k))]. \end{aligned}$$

Adding the terms for $k = \{1, \dots, K\}$ and dividing by K we get:

$$\begin{aligned} & \frac{1}{\epsilon} E \left[\sum_{l \in \mathcal{L}} w_l \mu_l^* - w_l \frac{1}{K} \sum_{k=1}^K I_l^*(a(k), c(k), d(k)) \right] \\ & \leq B_1 - \frac{E[V(d(K+1))]}{K} + \frac{E[V(d(1))]}{K} \\ & \leq B_1 + \frac{E[V(d(1))]}{K} \end{aligned} \quad (24)$$

where (24) follows from the fact that the Lyapunov function V is non-negative. Assuming $E[V(d(1))] < \infty$ we get the following limit expression:

$$\limsup_{K \rightarrow \infty} E \left[\sum_{l \in \mathcal{L}} w_l \mu_l^* - w_l \frac{1}{K} \sum_{k=1}^K I_l^*(a(k), c(k), d(k)) \right] \leq B\epsilon$$

where $B = B_1$. ■

APPENDIX C PROOF OF THEOREM 2

The proof will use dynamic programming arguments: we will first note that if there is only one time slot remaining, the optimal decision is to schedule the link with the highest weight among the links that have a packet that remains to be transmitted, and then using induction we will prove that the best decision in any time slot is to schedule a backlogged link with the highest weight.

For simplicity in notation, define $\pi_l = \frac{1}{\epsilon} w_l + d_l$ for all $l \in \mathcal{L}$. Also define

$$U_\rho(\mathcal{L}, j) = \sum_{l \in \mathcal{L}} \sum_c \sum_{t=T-j+1}^T \pi_l c_{lt} s_{lt}(\rho, a, c) Pr(c)$$

to be the expected utility of policy ρ when there are j time slots remaining and the set of links that have a packet that needs to be transmitted is given by \mathcal{L} . Furthermore, define

$$U_\rho(\mathcal{L}, 0) = 0.$$

Finally, we will denote the greedy policy by g .

If there is only one time slot remaining, the optimal decision is to schedule one of the links that has the largest weight $\pi_l \bar{c}_l$ among the links that have a waiting packet, since this maximizes the expected utility. So the optimal decision is to use the greedy scheduler in the last time slot. Using induction, we assume that when there are j time slots remaining, it is optimal to use the greedy scheduler. We will prove that if there are $j+1$ time slots remaining then it is also optimal to use the greedy scheduler.

When there are $j+1$ time slots, we need to determine which link to schedule in the first slot, and then use the greedy

scheduler for the remaining j time slots. Assume that the set of links that have a packet waiting to be transmitted is given by \mathcal{L} . If we schedule the link with the largest weight $\pi_l \bar{c}_l$, then the expected utility is given by

$$U_g(\mathcal{L}, j+1) = \pi_{l^*} \bar{c}_{l^*} + (1 - \bar{c}_{l^*}) U_g(\mathcal{L}, j) + \bar{c}_{l^*} U_g(\mathcal{L} \setminus \{l^*\}, j)$$

where

$$l^* \in \arg \max_{l \in \mathcal{L}} \pi_l \bar{c}_l.$$

If we decide to schedule link

$$\tilde{l} \notin \arg \max_{l \in \mathcal{L}} \pi_l \bar{c}_l,$$

then the expected utility is given by

$$\begin{aligned} U_{\tilde{\rho}}(\mathcal{L}, j+1) &= \pi_{\tilde{l}} \bar{c}_{\tilde{l}} + (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L}, j) + \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{\tilde{l}\}, j) \\ &= \pi_{\tilde{l}} \bar{c}_{\tilde{l}} + \pi_{l^*} \bar{c}_{l^*} + (1 - \bar{c}_{l^*}) (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L}, j-1) \\ &\quad + (1 - \bar{c}_{l^*}) \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{\tilde{l}\}, j-1) \\ &\quad + \bar{c}_{l^*} (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L} \setminus \{l^*\}, j-1) \\ &\quad + \bar{c}_{l^*} \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{l^*, \tilde{l}\}, j-1) \\ &= \pi_{l^*} \bar{c}_{l^*} + (1 - \bar{c}_{l^*}) [\pi_{\tilde{l}} \bar{c}_{\tilde{l}} + (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L}, j-1) \\ &\quad + \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{\tilde{l}\}, j-1)] \\ &\quad + \bar{c}_{l^*} [\pi_{\tilde{l}} \bar{c}_{\tilde{l}} + (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L} \setminus \{l^*\}, j-1) \\ &\quad + \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{l^*, \tilde{l}\}, j-1)]. \end{aligned}$$

So in order to prove that $U_g(\mathcal{L}, j+1) \geq U_{\tilde{\rho}}(\mathcal{L}, j+1)$ it suffices to show that

$$U_g(\mathcal{L}, j) \geq \pi_{\tilde{l}} \bar{c}_{\tilde{l}} + (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L}, j-1) + \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{\tilde{l}\}, j-1) \quad (25)$$

and

$$U_g(\mathcal{L} \setminus \{l^*\}, j) \geq \pi_{\tilde{l}} \bar{c}_{\tilde{l}} + (1 - \bar{c}_{\tilde{l}}) U_g(\mathcal{L} \setminus \{l^*\}, j-1) + \bar{c}_{\tilde{l}} U_g(\mathcal{L} \setminus \{l^*, \tilde{l}\}, j-1). \quad (26)$$

From the assumption that the greedy scheduler is optimal when there are j slots remaining, it is clear that (25) and (26) are true, so the greedy scheduler is indeed optimal when there are $j+1$ slots remaining. ■

REFERENCES

- [1] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," in *IEEE INFOCOM*, vol. 3, Miami, FL, USA, Mar. 13–17, 2005, pp. 1794–1803.
- [2] X. Lin and N. B. Shroff, "Joint rate control and scheduling in multihop wireless networks," in *43rd IEEE Conference on Decision and Control (CDC)*, vol. 2, Atlantis, Paradise Island, Bahamas, Dec. 14–17, 2004, pp. 1484–1489.
- [3] M. J. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *IEEE INFOCOM*, vol. 3, Miami, FL, USA, Mar. 13–17, 2005, pp. 1723–1734.
- [4] A. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, no. 4, pp. 401–457, Aug. 2005.
- [5] A. Eryilmaz and R. Srikant, "Joint congestion control, routing and mac for stability and fairness in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1514–1524, Aug. 2006.

- [6] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, "Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks," in *IEEE INFOCOM*, Barcelona, Catalunya, Spain, Apr. 23–29, 2006.
- [7] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [8] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1452–1463, Jun. 2006.
- [9] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Foundation and Trends in Networking*, vol. 1, no. 1, pp. 1–144, 2006.
- [10] S. Shakkottai and R. Srikant, "Scheduling real-time traffic with deadlines over a wireless channel," *Wireless Networks*, vol. 8, no. 1, pp. 13–26, Jan. 2002.
- [11] V. Raghunathan, V. Borkar, M. Cao, and P. R. Kumar, "Index policies for real-time multicast scheduling for wireless broadcast systems," in *IEEE INFOCOM*, Phoenix, AZ, USA, Apr. 13–18, 2008, pp. 1570–1578.
- [12] A. Dua and N. Bambos, "Downlink wireless packet scheduling with deadlines," *IEEE Trans. Mobile Comput.*, vol. 6, no. 12, pp. 1410–1425, Dec. 2007.
- [13] Q. Liu, X. Wang, and G. B. Giannakis, "A cross-layer scheduling algorithm with QoS support in wireless networks," *IEEE Trans. Veh. Technol.*, vol. 55, no. 3, pp. 839–847, May 2006.
- [14] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," in *IEEE INFOCOM*, Rio de Janeiro, Brazil, Apr. 19–25, 2009, pp. 486–494.
- [15] I.-H. Hou and P. R. Kumar, "Admission control and scheduling for QoS guarantees for variable-bit-rate applications on wireless channels," in *10th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, New Orleans, LA, USA, May 18–21, 2009, pp. 175–184.
- [16] —, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *IEEE INFOCOM*, San Diego, CA, USA, Mar. 15–19, 2010.
- [17] J. J. Jaramillo and R. Srikant, "Optimal scheduling for fair resource allocation in ad hoc networks with elastic and inelastic traffic," in *IEEE INFOCOM*, San Diego, CA, USA, Mar. 15–19, 2010.
- [18] D. G. Luenberger, *Linear and Nonlinear Programming*, 2nd ed. Norwell, MA: Kluwer Academic Publishers, 2003.
- [19] S. Boyd and L. Vandenberghe, *Convex Optimization*, 1st ed. New York, NY: Cambridge University Press, Mar. 2004.